

This Page Is Inserted by IFW Operations
and is not a part of the Official Record

BEST AVAILABLE IMAGES

Defective images within this document are accurate representations of the original documents submitted by the applicant.

Defects in the images may include (but are not limited to):

- BLACK BORDERS
- TEXT CUT OFF AT TOP, BOTTOM OR SIDES
- FADED TEXT
- ILLEGIBLE TEXT
- SKEWED/SLANTED IMAGES
- COLORED PHOTOS
- BLACK OR VERY BLACK AND WHITE DARK PHOTOS
- GRAY SCALE DOCUMENTS

IMAGES ARE BEST AVAILABLE COPY.

**As rescanning documents *will not* correct images,
please do not report the images to the
Image Problem Mailbox.**

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 462 917 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention
of the grant of the patent:
01.09.1999 Bulletin 1999/35

(51) Int Cl. 6: G06F 11/10

(21) Application number: 91480082.6

(22) Date of filing: 22.05.1991

(54) Method and apparatus for recovering parity protected data

Verfahren und Anordnung zur Wiederherstellung von durch Parität geschützten Daten

Procédé et dispositif de récupération de données protégées par parité

(84) Designated Contracting States:
DE FR GB IT

(30) Priority: 21.06.1990 US 542216

(43) Date of publication of application:
27.12.1991 Bulletin 1991/52

(73) Proprietor: International Business Machines
Corporation
Armonk, N.Y. 10504 (US)

(72) Inventors:

- Bond, Milton Fredrick
Rochester, Minnesota 55901 (US)
- Clark, Brian Eldridge
Rochester, Minnesota 55904 (US)
- McRoberts, Raymond Spencer
Rochester, Minnesota 55901 (US)

(74) Representative: de Pena, Alain et al
Compagnie IBM France
Département de Propriété Intellectuelle
06610 La Gaude (FR)

(56) References cited:
EP-A- 0 156 724 EP-A- 0 249 091

- Computer Architecture News vol. 17, no. 5,
September 1989, New York, USA pages 24 - 31;
R. H. Katz et al.: "A project on high performance
I/O subsystems"
- BYTE. vol. 15, no. 13, December 1990, ST
PETERBOROUGH US pages 337 - 339; M. H.
Anderson: "Strength [and safety] in numbers"
- COMPCON Spring '89, 34th IEEE Computer
Society International Conference, San
Francisco, CA, USA; Patterson et
al.: "Introduction to Redundant Arrays of
Inexpensive Disks (RAID)"

EP 0 462 917 B1

Note: Within nine months from the publication of the mention of the grant of the European patent, any person may give notice to the European Patent Office of opposition to the European patent granted. Notice of opposition shall be filed in a written reasoned statement. It shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

[0001] The present invention relates to maintaining parity information on computer data storage devices and in particular to maintaining availability of a computer system when reconstructing data from a failed storage device.

[0002] The extensive data storage needs of modern computer systems require large capacity mass data storage devices. A common storage device is the magnetic disk drive, a complex piece of machinery containing many parts which are susceptible to failure. A typical computer system will contain several such units. As users increase their need for data storage, systems are configured with larger numbers of storage units. The failure of a single storage unit can be a very disruptive event for the system. Many systems are unable to operate until the defective unit is repaired or replaced, and the lost data restored. An increased number of storage units increases the probability that any one unit will fail, leading to system failure. At the same time, computer users are relying more and more on the consistent availability of their systems. It therefore becomes essential to find improved methods of reconstructing data contained on a failing storage unit, and sustaining system operations in the presence of a storage unit failure.

[0003] One method of addressing these problems is known as "mirroring". This method involves maintaining a duplicate set of storage devices, which contains the same data as the original. The duplicate set is available to assume the task of providing data to the system should any unit in the original set fail. Although very effective, this is a very expensive method of resolving with the problem since a customer must pay for twice as many storage devices.

[0004] A less expensive alternative is the use of parity blocks. Parity blocks are records formed from the Exclusive-OR of all data records stored at a particular location on different storage units. In other words, each bit in a block of data at a particular location on a storage unit is Exclusive-ORed with every other bit at that same location in each storage unit in a group of units to produce a block of parity bits; the parity block is then stored at the same location on another storage unit. If any storage unit in the group fails, the data contained at any location on the failing unit can be regenerated by taking the Exclusive-OR of the data blocks at the same location on the remaining devices and their corresponding parity block.

[0005] U.S. Pat. No. 4,092,732 (Ouchi et al.) describes a parity block method. In said patent, a single storage unit is used to store parity information for a group of storage devices. A read and a write on the storage unit containing parity blocks occurs each time a record is changed on any of the storage units in the group covered by the parity record. Thus, the storage unit with the parity records becomes a bottleneck to storage operations. EP-A-0 249 091 (Clark et al.) improves

upon storage of parity information by distributing parity blocks substantially equally among a set of storage units. N storage units in a set are divided into a multiple of equally sized address blocks, each containing a plurality of records. Blocks from each storage unit having the same address ranges form a stripe of blocks. Each stripe has a block on one storage device containing parity for the remaining blocks of the stripe. The parity blocks for different stripes are distributed among the different storage units in a round robin manner.

[0006] The use of parity records as described in the Ouchi and Clark patents substantially reduces the cost of protecting data when compared to mirroring. However, while Ouchi and Clark teach a data recovery or protection means, they do not provide a means to keep a system operational to a user during data reconstruction. Normal operations are interrupted while a memory controller is powered down to permit a repair or replacement of the failed storage device, followed by a reconstruction of the data. Since this prior art relies exclusively on software for data reconstruction, the system can be disabled for a considerable time.

[0007] Prior art does not teach dynamic system recovery and continued operation without the use of duplicate or standby storage units. Mirroring requires a doubling of the number of storage units. A less extreme approach is the use of one or more standby units, i.e., additional spare disk drives which can be brought on line in the event any unit in the original set fails. Although this does not entail the cost of a fully mirrored system, it still requires additional storage units which otherwise serve no useful function.

[0008] An article of Randy H. Katz 'A project on High Performance I/O Subsystems' in '88345 Computer Architecture News (September 1989) discloses the Disk Arrays providing High reliability. The array of disks is organized into separate groups for recovery purpose. High reliability is reached at very low cost by just adding one extra disk per group to hold the parity calculated horizontally across the group. Redundancy, spread of data and parity bits over the storage units are the characteristics of the N+1 RAID solution discussed in the article 'A case for Redundant Arrays of Inexpensive Disks' of ACM SIGMO Conference Chicago, Illinois, June 1-3, 1988. N+1 RAID is also the level 5 RAID as exposed in a previous article of the same authors. In the article the mirrored RAID and the N+1 RAID solutions are compared in terms of cost and performance. The parity of one group of disks is calculated on a per-bit basis any

single disk failure can be corrected by reading the rest of the disks in the group to determine what bit value on the failed disk would give the proper parity. A bottleneck of the parity disk is avoided by spreading the parity over several disks (page 113 5. N+1 RAID paragraph). Even if the articles provides a solution with low cost disks and an interesting ratio cost/performance (in terms of reliability and data access time), none of these documents deals with the problems of dynamic reconstruction of data.

ta from a failure; these problems concern the storage space used for reconstruction of data and the availability of the disk array for the computer system using it.

[0009] It is therefore an object of the present invention to provide an enhanced method and apparatus for recovering from data loss in a computer system having multiple data storage units.

[0010] It is a further object of this invention to provide an enhanced method and apparatus whereby a computer system having multiple data storage units may continue to operate if one of the data storage units fails.

[0011] Another object of this invention is to reduce the cost of protecting data in a data processing system having multiple protected storage units.

[0012] A still further object of this invention is to increase the performance of a computer system having multiple data storage units when one of the data storage units fails and the system must reconstruct the data contained on the failed unit.

[0013] A storage controller services a plurality of data storage units. A storage management mechanism resident on the controller maintains parity records on the storage units it services. Data and parity blocks are organized as described in the patent to Clark et al. In the event of a storage unit failure, the system continues to operate. The storage management mechanism reconstructs data that was on the failed unit as attempts are made to access that data, and stores it in the parity block areas of the remaining storage units.

[0014] The storage management mechanism includes a status map indicating, for each data block, the location of the corresponding parity block, and the status of the data block. If a storage unit fails, the storage management mechanism is placed in a failure operating mode. While in failure operating mode, the storage management mechanism checks the status map before accessing data on the failed storage unit. If the data has not yet been reconstructed, storage management must first reconstruct the data in that block of storage by successively reading and accumulating an Exclusive-OR (XOR) of the same blocks on all storage units in the parity group, including the parity block. The block of data resulting from this Exclusive-OR is the reconstructed data, which is then stored in the location of the parity block. The status map is then updated to indicate that the block has been reconstructed. Once the data has been reconstructed, it is only necessary to read from or write to the former parity block directly. In the same manner, storage management will reconstruct the data from a block of storage on the failed unit before writing to any other block on the same stripe (on a non-failed unit). This is required because the write operation to any block on the stripe will alter parity, making it impossible to later reconstruct the block of data on the failed unit. Thus, upon failure of a storage unit, system performance is initially degraded as read and write operations cause storage management to reconstruct data. As data is rebuilt, performance quickly improves.

[0015] In the preferred embodiment, the storage units are organized and parity information is generated and stored as described in the Clark et al. patent. Reconstructed data is stored in locations where parity data is normally stored for the stripe on which the lost data resided. There is no need to power down the storage controller or any other part of the system, repair the failed storage unit, and then reconstruct the lost data. In this preferred embodiment, the data are recovered and stored while a computer system using this storage management mechanism remains completely available to a user. The storage units operate without parity protection until the failed unit is repaired or replaced. This embodiment achieves continuous operation and single-level failure protection at very little additional cost.

[0016] In a first alternate embodiment, spare areas of storage in each non-failing storage unit are allocated to the reconstructed data. The total of these spare areas constitute a virtual spare storage unit. As data is reconstructed, it is placed in the virtual spare unit, and parity is maintained in the normal fashion. This alternative achieves an additional level of failure protection, because parity data continues to be maintained after a single storage unit failure. However, it may impose a need for additional storage space for the spare areas, or cause degraded performance if these spare areas are normally used for other purposes, such as temporary data storage.

[0017] In a second alternate embodiment, the storage management mechanism resides in the host system's operating software, but otherwise performs the same functions as a storage management mechanism residing on a storage controller. This embodiment will generally be slower than the preferred embodiment, but may reduce the cost of the storage controller.

Fig. 1 is a block diagram of a system incorporating the components of the preferred embodiment of this invention;

Fig. 2 is a diagram of a status map;

Fig. 3 is a flow diagram of the steps involved in a read operation during normal operating mode;

Fig. 4 is a flow diagram of the steps involved in transferring data to be written from the host to the storage controller;

Fig. 5 is a flow diagram of the steps involved in writing data to a storage device in normal operating mode;

Fig. 6 is a flow diagram of steps involved in read operations following a storage device failure.

Fig. 7 is a flow diagram of the steps involved in writing data to a storage device when a storage device

has failed;

Fig. 8 is a block diagram of a system incorporating the components according to an alternative embodiment of this invention.

[0018] A block diagram of the major components of computer system 100 of the preferred embodiment of the present invention is shown in Figure 1. A host system 101, communicates over a bus 102 with a storage controller 103. Controller 103 comprises a programmed processor 104, non-volatile RAM 105, Exclusive-OR hardware 108, and cache memory (RAM) 109. Non-volatile RAM 105 contains a status map 106 and table of contents 107. Controller 103 controls the operation of storage units 121-124. In the preferred embodiment, units 121-124 are rotating magnetic disk storage units. While four storage units are shown in Fig. 1, it should be understood that the actual number of units attached to controller 103 is variable. It should also be understood that more than one controller 103 may be attached to host system 101. In the preferred embodiment, computer system 100 is an IBM AS/400 computer system, although any computer system could be used.

[0019] The storage area of each storage unit is divided into blocks 131-138. In the preferred embodiment, all storage units have identical storage capacity, and all parity protected blocks the same size. While it would be possible to employ this invention in configurations of varying sized storage units or varying sized blocks, the preferred embodiment simplifies the control mechanism.

[0020] The set of all blocks located at the same location on the several storage units constitutes a stripe. In Fig. 1, storage blocks 131-134 constitute a first stripe, and blocks 135-138 constitute a second stripe. One of the blocks in each stripe is designated the parity block. Parity blocks 131,136 are shown shaded in Fig. 1. The remaining unshaded blocks 132-135,137-138 are data storage blocks for storing data. The parity block for the first stripe, consisting of blocks 131-134, is block 131. The parity block contains the Exclusive-OR of data in the remaining blocks on the same stripe.

[0021] In the preferred embodiment, parity blocks are distributed across the different storage units in a round robin manner, as shown in Fig. 1. Because with every write operation the system must not only update the block containing the data written to, but also the parity block for the same stripe, parity blocks are usually modified more frequently than data blocks. Distributing parity blocks among different storage units will in most cases improve performance by distributing the access workload. However, such distribution is not necessary to practicing this invention, and in an alternate embodiment it would be possible to place all parity blocks on a single storage unit.

[0022] In the preferred embodiment, one block of every stripe is dedicated to parity information. As an alter-

native embodiment, one of the stripes contains no parity protection. This stripe is reserved for temporary data which does not require protection. Fig. 8 shows this alternate embodiment in the stripe consisting of blocks 811-814. Because it is extra storage space not a part of the parity data protection scheme, this block may be of any arbitrary size.

[0023] The allocation of storage area into stripes as described above, each containing blocks of data and a parity block, is the same as that described in U.S. Patent 4,761,785 to Clark, et al., which is incorporated by reference.

[0024] Storage controller 103 includes programmed processor 104 executing a storage management program. The operation of the storage management program is described below. Controller 103 also includes hardware Exclusive-OR circuitry 108, for computing the Exclusive-OR of data in non-volatile RAM 105 or cache RAM 109. In an alternative embodiment, the Exclusive-OR operations could be performed by processor 104, but special hardware for this purpose will improve performance.

[0025] Non-volatile RAM 105 is used by controller 103 as a temporary queueing area for data waiting to be physically written to a storage unit. In addition to this temporary data, status map 106 and table of contents 107 are stored in non-volatile RAM 105. Table of contents 107 contains a mapping of the data waiting to be written to the location on which it is stored in the storage unit.

[0026] Status map 106 is used to identify the location of the corresponding parity block for each data block, and the status of each block of data during failure recovery mode. Status map 106 is shown in detail in Fig. 2. It contains a separate table of status map entries for each storage unit. Each status map entry 201 contains the location 202 of a block of data on the storage unit, a status bit 203 indicating whether or not the data needs to be recovered when operating in failure mode, and the location of the corresponding parity block 204.

[0027] Referring again to Fig. 1, cache memory 109 is a volatile random access memory that is used to store data read from a storage unit. It serves as a buffer when transferring data from a storage unit to host system 101 in a read operation. In addition, data is saved in cache 109 in response to indications from the host system 101 that the data has a high probability of modification and re-writing. Because unmodified data must be exclusive-ORed with modified data to update the corresponding parity data, saving read data in cache 109 can eliminate the need to read it again immediately before a write operation. Cache 109 exists only to improve performance. In an alternative embodiment, it would be possible to practice this invention without it. Cache 109 is identified as a volatile RAM because it is not necessary to the integrity of the system that data read from storage be preserved in non-volatile memory. However, the cache could be implemented as part of the non-volatile mem-

ory 105. Depending on the relative cost and size of memory modules, such an approach may be desirable. [0028] The function of the system in conjunction with the hardware and software features necessary to this invention is described below. The system has two operating modes: normal and failure mode. The system operates in normal mode when all disk storage devices are functioning properly. When one storage device fails, the mode of operation changes to failure mode, but the system continues to operate.

[0029] A READ operation in normal mode is shown in Fig. 3. The READ operation is performed by accepting a READ command from the host at step 301, and determining whether the data requested exists in non-volatile RAM 105 or cache 109 at step 302. If so, the data in non-volatile RAM or cache is sent directly to the host at step 304. Otherwise, data is first read from the appropriate storage unit into the cache 109 at step 303, and from there transferred to the host system at step 304. The cache 109 also improves performance during WRITE operations. If the original version of data to be updated is already in cache 109 when a WRITE operation is processed, it is not necessary to read the data again in order to update parity, thus improving system performance. The contents of cache 109 are managed using any of various cache management techniques known in the art.

[0030] A WRITE operation is performed by two asynchronous tasks running in the storage controller's processor 104. One task communicates with the host via bus 102, and is shown in Fig. 4. The WRITE operation begins when it accepts a WRITE command from the host at step 401. It then checks table of contents 107 to determine whether sufficient space is available in non-volatile RAM 105 to store the data to be written to storage in step 402 (Note that space available includes space used by back-level versions of the data to be written, as well as unused space). If space is not available, controller 103 can not receive data from the host, and must wait for space to become available at step 403 (i.e., it must wait for data already in non-volatile RAM 105 to be written to storage 121-124). When space becomes available in non-volatile RAM 105, data is copied from host 101 into non-volatile RAM 105, and table of contents 107 is updated at step 404. Processor 104 then issues an operation complete message to the host at step 405. Upon receipt of the operation complete message, the host is free to continue processing as if the data were actually written to storage 121-124, although in fact the data may wait awhile in non-volatile RAM 105. From the host's perspective, the operation will appear to be complete.

[0031] The second asynchronous task writes data from non-volatile RAM 105 to a storage unit. A flow diagram of this task in normal mode is shown in Fig. 5. The task selects a WRITE operation from among those queued in non-volatile RAM at step 501. The selection criteria are not a part of this invention, and could be, e.

g., First-in-first-out, Last-in-first-out, or some other criteria based on system performance and other considerations. When the WRITE operation is performed, parity must be updated. By taking the Exclusive-OR of the new write data with the old data, it is possible to obtain a bit map of those bits being changed by the WRITE operation. Exclusive-ORing this bit map with the existing parity data produces the updated parity data. Therefore, before writing to storage, the task first checks whether the old data exists in the cache 109 in unmodified form at step 502. If not, it is read into the cache from storage at step 503. This old data in the cache is then Exclusive-ORed with the new data in non-volatile RAM to produce the bit map of changed data at step 504. The bit map is saved temporarily in non-volatile RAM 105 while the new data is written to one of the storage devices 121-124. The old parity data is then read into the cache (if not already there) at steps 506, 507, and Exclusive-ORed with the bit map to produce the new parity data at step 508. This new parity data is written to one of the storage devices 121-124 and the table of contents is updated at step 509, completing the WRITE operation. [0032] When a storage unit failure is detected, the system begins operating in failure mode. The failure of a storage unit means failure to function, i.e., to access data. Such a failure is not necessarily caused by a breakdown of the unit itself. For example, the unit could be powered off, or a data cable may be disconnected. From the perspective of the system, any such failure, whatever the cause, is a failure of the storage unit. Detection mechanisms which detect such failures are known in the art. Common mechanisms include a timeout after not receiving a response, and continued high error rates in received data. [0033] Figure 6 illustrates the READ operation when the system is operating in failure mode. As in the case of normal mode READ operations, when a READ is accepted from the host at step 601, the controller first checks its non-volatile RAM 105 and its volatile cache 109 for the desired data at step 602. If the data exists in non-volatile RAM or cache, the data is transferred to the host via system bus 102. If the data is not in non-volatile RAM or cache, and resides on a storage device which has not failed (step 603), the data is read into the cache from the storage device in the normal manner at step 604. If the data resides on a failed storage unit, the controller checks the status map entry 201 in status map 106 for the location in storage of the desired data at step 605. The status map entry will indicate whether the data has been recovered, i.e., whether it has been reconstructed by exclusive-ORing and stored at some alternate location. If the status map indicates that the data has not been recovered (step 605) the controller successively reads the corresponding locations on all storage units except the failing one at step 608. Each block of data read is XORed by the XOR hardware 108 with the accumulated XOR results of the previously read blocks. The final XOR results constitute the reconstruct-

ed data of the failed device. This reconstructed data is written to the parity block corresponding to this block of data at step 609. The location of this block is stored in a parity block address field 204 of the status map 108. After writing the recovered data to the parity block location, status map 108 is updated at step 610 by changing the status bit 203 of each block in the same stripe to a '1' to indicate that the data has been recovered. The reconstructed data is sent to the host at step 611. If the status bit 203 originally contained a '1', indicating that data had been recovered, the controller would obtain the location of the former parity block area (where recovered data is stored) from the status map at step 606, and read the data from this location directly into the cache at step 607. By this device, it is only necessary to read all disk storage units once to recover any particular block of data. Once recovered, the physical storage location of that data is effectively relocated to the location that was formerly used for parity storage, and any subsequent reads of that block need only read the one storage unit.

[0034] Figure 7 illustrates the write to storage operation when the system is operating in failure mode. As with the normal mode WRITE, a host communications task shown in Fig. 4 receives data to be written from the host via bus 102. The write to storage task selects a write operation from the queue in non-volatile RAM 105 at step 701. The controller determines whether the data is to be written to a failed unit (step 702) and checks the status map (steps 703, 709). If the data is to be written to a failing unit, and the data in the block has not yet been recovered, the block must be recovered before any write operations are possible. Recovery follows the same steps described above for a READ operation. Each block in the same stripe of blocks (including the parity block) is read in turn, and its contents Exclusive-ORed with the cumulative Exclusive-OR of the previously read blocks at step 704. The result, which is the reconstructed data, is written to the location used for the parity block at step 705. Once the recovery of the entire block is complete, the new data (which would typically encompass only a portion of the block) is written over the recovered data in the former parity location at step 706, and the status map updated to indicate that the block has been recovered at step 707. If data is to be written to a failing unit, but the data has already been recovered, it is written directly to the former parity location, now used for storage of recovered data, at step 708.

[0035] If data is being written to a non-failing unit when operating in failure mode, the controller checks the status map at step 709. If the status is '1', indicating that the block of data in the same stripe on the failing unit has already been recovered, the WRITE data is written directly to the non-failing storage unit at step 710. If the status is '0', data can not be directly written to the non-failing unit, because such an operation would alter parity, making it impossible to later reconstruct the cor-

responding data in the failed unit. Accordingly, in the preferred embodiment, the controller will first recover the block of data in the same stripe on the failing unit. As shown in Fig. 7, the block of data in the failing unit is first reconstructed by Exclusive-ORing at step 711, and saved in the parity block location at step 712, following the steps described above. The WRITE data is then written to its storage unit at step 713, and the status map is updated at step 714. Note that if the parity block for the stripe containing the data to be written is on the failing unit, no reconstruction is necessary, since parity will be lost anyway. Therefore, the status for all blocks on this stripe is set to 1 when the storage unit failure is detected. The effect will be to cause data on this stripe to be directly written to storage as if the corresponding block on the failing unit had already been recovered. For example, referring to Fig. 1, if storage unit 121 fails, the controller will immediately set the status of blocks 132-134 to '1', so that WRITE operations to these blocks can proceed directly. In an alternative embodiment, if the WRITE operation is to a non-failing unit, and the corresponding block on the failing unit has not been recovered, it would be possible to follow the same steps used for a normal mode WRITE operation to update the parity block, preserving the ability to reconstruct the failing unit's data later if a READ or WRITE of the data on the failed unit is requested.

[0036] In the preferred embodiment, parity blocks are used to store reconstructed data, with the result that the system runs without parity protection after a single storage unit failure. An alternative embodiment is possible where a sufficiently large spare storage stripe or stripes is reserved on the storage units, as shown in Fig. 8. This spare storage stripe might contain temporary data which does not require parity protection and which can be overwritten if the need arises, or it might contain no data at all. In this alternative embodiment, reconstructed data is relocated to a block of a spare storage stripe 811-814 instead of the parity block. This alternative is only possible where sufficient spare storage exists to accommodate the non-spare contents of the failed unit. It would also have the consequence of reducing the amount of temporary storage available to the system, possibly degrading performance or reducing the number of users the system can service. In this alternative embodiment, normal mode READ and WRITE operations are performed in exactly the same manner as in the preferred embodiment. When operating in failure mode, the status map is checked, and the data reconstructed as needed, in the manner described above. However, instead of writing the reconstructed data to the parity block, it is written to a block in spare storage. Another field is required in status map 106 to record the new location of the data which was contained on the failed unit. In addition, with any WRITE operation parity is updated in the same manner as a WRITE operation in normal mode. This is done after any reconstruction of data on the failed unit.

[0037] In another alternative embodiment, parity protection and mirroring are combined on the same system. Some of the data contained on the storage units is protected by the parity protection mechanism described herein, while other data is mirrored. In the event of a storage unit failure, the parity protected data is reconstructed and stored as described above, while the mirrored data is accessed from the storage unit containing the mirrored copy.

[0038] Although a specific embodiment of the invention has been disclosed along with certain alternatives, it will be recognized by those skilled in the art that additional variations in form and detail may be made. In particular, while the disclosed preferred embodiment employs magnetic disk storage units, the invention is applicable to other storage device technologies having erasable, read/write characteristics.

Claims

1. A method of operating a computer system having a set of storage blocks, said set comprising a plurality of data storage blocks for containing data and a parity protection storage block for containing the parity protection information of the data stored in said data storage blocks, each of said storage blocks being contained on a respective data storage unit, said method comprising the steps of:

reconstructing data contained in a storage block, while the data storage unit containing said storage block is failing, from the remaining storage blocks in the set;

characterized by

storing data reconstructed by said reconstructing step on one of said data storage units other than said failing data storage unit.

2. The method of operating a computer system of claim 1, wherein said reconstructing data step reconstructs data when attempts are made to access said data.

3. The method of operating a computer system of claim 1 or 2, wherein said storing data step stores the reconstructed data in said parity protection storage block.

4. The method of operating a computer system of claim 3, wherein said reconstructing data step reconstructs data when attempts are made to access said data.

5. The method of operating a computer system of any one of claims 1-4, wherein said data storage units

5 contain a spare storage block, and said storing data step stores the reconstructed data in said spare storage block. A storage apparatus for a computer system, comprising:

at least three data storage units;

10 at least one set of storage blocks, each set comprising a plurality of data storage blocks for containing data and a parity protection storage block for containing the parity protection of the data stored in said data storage blocks, each of said storage blocks being contained on a respective data storage unit;

15 means for reconstructing the data contained in one of said data storage blocks, while the data storage unit containing said block is failing, from the remaining storage blocks in the set; and

20 means for storing said reconstructed data on one of said data storage units.

25 6. A storage apparatus for a computer system, comprising

at least three data storage units;

30 at least one set of storage blocks, each set comprising a plurality of data storage blocks for containing data and a parity protection storage block for containing the parity protection of the data stored in said data storage blocks, each of said storage blocks being contained on a respective data storage unit;

35 means for reconstructing the data contained in one of said data storage blocks, while the data storage unit containing said block is failing, from the remaining storage blocks in the set;

characterized by

40 means for storing said reconstructed data on one of said data storage units other than said failing data storage unit.

45 7. The storage apparatus for the computer system of claim 6, wherein said means for storing said reconstructed data stores said data in said parity protection storage block.

50 8. The storage apparatus for the computer system of claim 7, wherein said means for reconstructing the data comprises a storage controller, said storage controller comprising:

a programmable processor executing a storage management program; and

a non-volatile random access memory.

9. The storage apparatus for the computer system of claim 7, wherein said data processing system comprises at least two of said set of storage blocks; and

wherein said parity protection blocks are distributed among said data storage units in a round robin manner.

10. The storage apparatus for the computer system of claim 7, wherein each of said data storage units is a rotating magnetic disk drive storage unit.

11. The storage apparatus for the computer system of anyone of claims 6-10, wherein each of said data storage units contains a spare storage block; and

wherein said means for storing said reconstructed data stores said data in one of said spare storage blocks.

Patentansprüche

1. Ein Verfahren zum Betreiben eines Rechnersystems mit einem Satz Speicherblöcken, wobei dieser Satz eine Vielzahl von Datenspeicherblöcken zum Aufnehmen der Daten und einen Paritätsschutz-Speicherblock zum Aufnehmen der Paritätsschutzinformationen der in den Datenspeicherblöcken gespeicherten Daten aufweist, wobei jeder der Speicherblöcke auf einer entsprechenden Datenspeichereinheit enthalten ist, und das Verfahren folgende Schritte aufweist:

Wiedergewinnen der in einem Speicherblock enthaltenen Daten bei Ausfall der Datenspeichereinheit, die diesen Speicherblock enthält, aus den restlichen Speicherblöcken in dem betreffenden Satz;

gekennzeichnet durch

das Speichern der durch diesen Wiedergewinnungsschritt rekonstruierten Daten auf einer dieser Datenspeichereinheiten, die nicht die ausgefallene Datenspeichereinheit ist.

2. Das Verfahren zum Betreiben eines Rechnersystems gemäß Anspruch 1, in dem der Datenwiedergewinnungsschritt die Daten rekonstruiert, wenn Versuche gemacht werden, auf diese Daten Zugriff zu nehmen.

3. Das Verfahren zum Betreiben eines Rechnersystems gemäß Anspruch 1 oder 2, in dem der Datenwiedergewinnungsschritt die rekonstruierten Daten im Paritätsschutz-Speicherblock speichert.

4. Das Verfahren zum Betreiben eines Rechnersystems gemäß Anspruch 3, in dem der Datenwiedergewinnungsschritt die Daten rekonstruiert, wenn Versuche gemacht werden, auf diese Daten Zugriff zu nehmen.

5. Das Verfahren zum Betreiben eines Rechnersystems gemäß einem beliebigen der Ansprüche 1-4, in dem die Datenspeichereinheiten einen Ersatzspeicherblock enthalten, und der Datenspeicherschritt die rekonstruierten Daten in dem Ersatzspeicherblock speichert. Eine Speichervorrichtung für ein Rechnersystem, enthaltend:

wenigstens drei Datenspeichereinheiten;

wenigstens einen Satz Speicherblöcke, wobei jeder Satz eine Vielzahl von Datenspeicherblöcken zur Aufnahme von Daten sowie einen Paritätsschutz-Speicherblock zur Aufnahme des Paritätsschutzes der in den Datenblöcken gespeicherten Daten, wobei jeder der Speicherblöcke in einer entsprechenden Datenspeichereinheit enthalten ist;

Mittel zum Rekonstruieren der in einem der Datenspeicherblöcke enthaltenen Daten bei Ausfall der die Blöcke enthaltenden Datenspeichereinheit aus den restlichen Speicherblöcken im Satz; und

Mittel zum Speichern der rekonstruierten Daten auf einer der Datenspeichereinheiten.

6. Eine Speichervorrichtung für ein Rechnersystem, enthaltend:

Wenigstens drei Datenspeichereinheiten;

wenigstens einen Satz Speicherblöcke, wobei jeder Satz eine Vielzahl von Datenspeicherblöcken zur Aufnahme von Daten sowie einen Paritätsschutz-Speicherblock zur Aufnahme des Paritätsschutzes der in den Datenblöcken gespeicherten Daten, wobei jeder der Speicherblöcke in einer entsprechenden Datenspeichereinheit enthalten ist;

Mittel zum Rekonstruieren der in einem der Datenspeicherblöcke enthaltenen Daten bei Ausfall der die Blöcke enthaltende Datenspeichereinheit aus den restlichen Speicherblöcken im Satz;

gekennzeichnet durch

Mittel zum Speichern der rekonstruierten Daten auf einer der Datenspeichereinheiten, die nicht die ausgefallene Datenspeichereinheit ist. 5

7. Die Speichervorrichtung für ein Rechnersystem gemäß Anspruch 6, in der das Mittel zum Speichern der rekonstruierten Daten diese Daten in den Prioritätsschutz-Speicherblock speichert. 10

8. Die Speichervorrichtung für ein Rechnersystem gemäß Anspruch 7, in der das Mittel zum Rekonstruieren der Daten einen Speicher-Controller umfaßt, wobei dieser Speicher-Controller enthält:

einen programmierbaren Prozessor, der ein Speicherverwaltungsprogramm abarbeitet; und 15

Einen nichtflüchtigen Direktzugriffsspeicher. 20

9. Die Speichervorrichtung für das Rechnersystem gemäß Anspruch 7, in der das Datenbearbeitungssystem mindestens zwei Speicherblöcke des Speicherblockatzes enthält; und

in dem die Paritätsschutzböcke unter den Datenspeichereinheiten in Zeitrasterfolge rundum verteilt sind. 25

10. Die Speichervorrichtung für ein Rechnersystem gemäß Anspruch 7, in der jede der Datenspeichereinheiten eine rotierende Magnetplattenlaufwerk-Speichereinheit ist.

11. Die Speichervorrichtung für ein Rechnersystem gemäß einem beliebigen der Ansprüche 6-10, in dem jede der Datenspeichereinheiten einen Ersatzspeicherblock enthält; und

in dem das Mittel zum Speichern der rekonstruierten Daten diese Daten in einem der Ersatzspeicherblöcke speichert. 30

Revendications

1. Procédé de mise en oeuvre d'un système d'ordinateur comportant un ensemble de blocs de mémorisation, ledit ensemble comprenant une pluralité de blocs de mémorisation de données destinés à contenir des données et un bloc de mémorisation de protection de parité destiné à contenir les informations de protection de parité des données mémorisées dans lesdits blocs de mémorisation de données, chacun desdits blocs de mémorisation étant contenu sur une unité de mémorisation de données 50

respective, ledit procédé comprenant les étapes consistant à :

reconstituer les données contenues dans un bloc de mémorisation, alors que l'unité de mémorisation de données contenant ledit bloc de mémorisation est défaillante, d'après les blocs de mémorisation restant dans l'ensemble,

caractérisé par

des données de mémorisation reconstituées par ladite étape de reconstitution sur l'une des unités de mémorisation de données autre que ladite unité de mémorisation de données défaillante.

2. Procédé de mise en oeuvre d'un système d'ordinateur selon la revendication 1, dans lequel ladite étape de reconstitution des données reconstitue des données lorsque des tentatives sont faites d'accéder auxdites données. 20

3. Procédé de mise en oeuvre d'un système d'ordinateur selon la revendication 1 ou 2, dans lequel ladite étape de mémorisation des données mémorise les données reconstituées dans ledit bloc de mémorisation de protection de parité. 25

4. Procédé de mise en oeuvre d'un système d'ordinateur selon la revendication 3, dans lequel ladite étape de reconstitution des données reconstitue les données lorsque des tentatives sont faites d'accéder auxdites données. 30

5. Procédé de mise en oeuvre d'un système d'ordinateur selon l'une quelconque des revendications 1 à 4, dans lequel lesdites unités de mémorisation de données contiennent un bloc de mémorisation de secours, et ladite étape de mémorisation de données mémorise les données reconstituées dans ledit bloc de mémorisation de secours. Dispositif de mémorisation destiné à un système d'ordinateur, comprenant :

au moins trois unités de mémorisation de données,

au moins un ensemble de blocs de mémorisation, chaque ensemble comprenant une pluralité de blocs de mémorisation de données destinés à contenir des données et un bloc de mémorisation de protection de parité destiné à contenir la protection de parité des données mémorisées dans lesdits blocs de mémorisation, chacun desdits blocs de mémorisation étant contenu sur une unité de mémorisation de données respective,

un moyen destiné à reconstituer les données contenues dans l'un desdits blocs de mémorisation de données, alors que l'unité de mémorisation de données contenant ledit bloc est défaillante, d'après les blocs de mémorisation restant dans l'ensemble, et

un moyen destiné à mémoriser lesdites données reconstituées sur l'une desdites unités de mémorisation de données.

6. Dispositif de mémorisation destiné à un système d'ordinateur, comprenant

au moins trois unités de mémorisation de données,

au moins un ensemble de blocs de mémorisation, chaque ensemble comprenant une pluralité de blocs de mémorisation de données destinés à contenir des données et un bloc de mémorisation de protection de parité destiné à contenir la protection de parité des données mémorisées dans lesdits blocs de mémorisation de données, chacun desdits blocs de mémorisation étant contenu sur une unité de mémorisation de données respective,

un moyen destiné à reconstituer les données contenues dans l'un desdits blocs de mémorisation de données, alors que l'unité de mémorisation de données contenant ledit bloc est défaillante, d'après les blocs de mémorisation restant dans l'ensemble, et

9. Dispositif de mémorisation destiné au système d'ordinateur selon la revendication 7, dans lequel ledit système de traitement de données comprend au moins deux blocs dudit ensemble de blocs de mémorisation,

dans lequel lesdits blocs de protection de parité sont répartis parmi lesdites unités de mémorisation de données suivant une séquence périodique.

10. Dispositif de mémorisation destiné au système d'ordinateur selon la revendication 7, dans lequel chacune desdites unités de mémorisation de données est une unité de mémorisation à lecteur de disque magnétique rotatif.

11. Dispositif de mémorisation destiné au système d'ordinateur selon l'une quelconque des revendications 6 à 10, dans lequel chacune desdites unités de mémorisation de données contient un bloc de mémorisation de secours, et

dans lequel ledit moyen destiné à mémoriser lesdites données reconstituées mémorise lesdites données dans l'un desdits blocs de mémorisation de secours.

caractérisé par

un moyen destiné à mémoriser lesdites données reconstituées sur l'une desdites unités de mémorisation de données, autre que ladite unité de mémorisation de données défaillante.

7. Dispositif de mémorisation destiné au système d'ordinateur selon la revendication 6, dans lequel ledit moyen destiné à mémoriser lesdites données reconstituées mémorise lesdites données dans lesdits blocs de mémorisation de protection de parité.

8. Dispositif de mémorisation destiné au système d'ordinateur selon la revendication 7, dans lequel ledit moyen destiné à reconstituer les données comprend un contrôleur de mémorisation, ledit contrôleur de mémorisation comprenant :

un processeur programmable exécutant un programme de gestion de mémorisation, et

une mémoire vive non volatile.

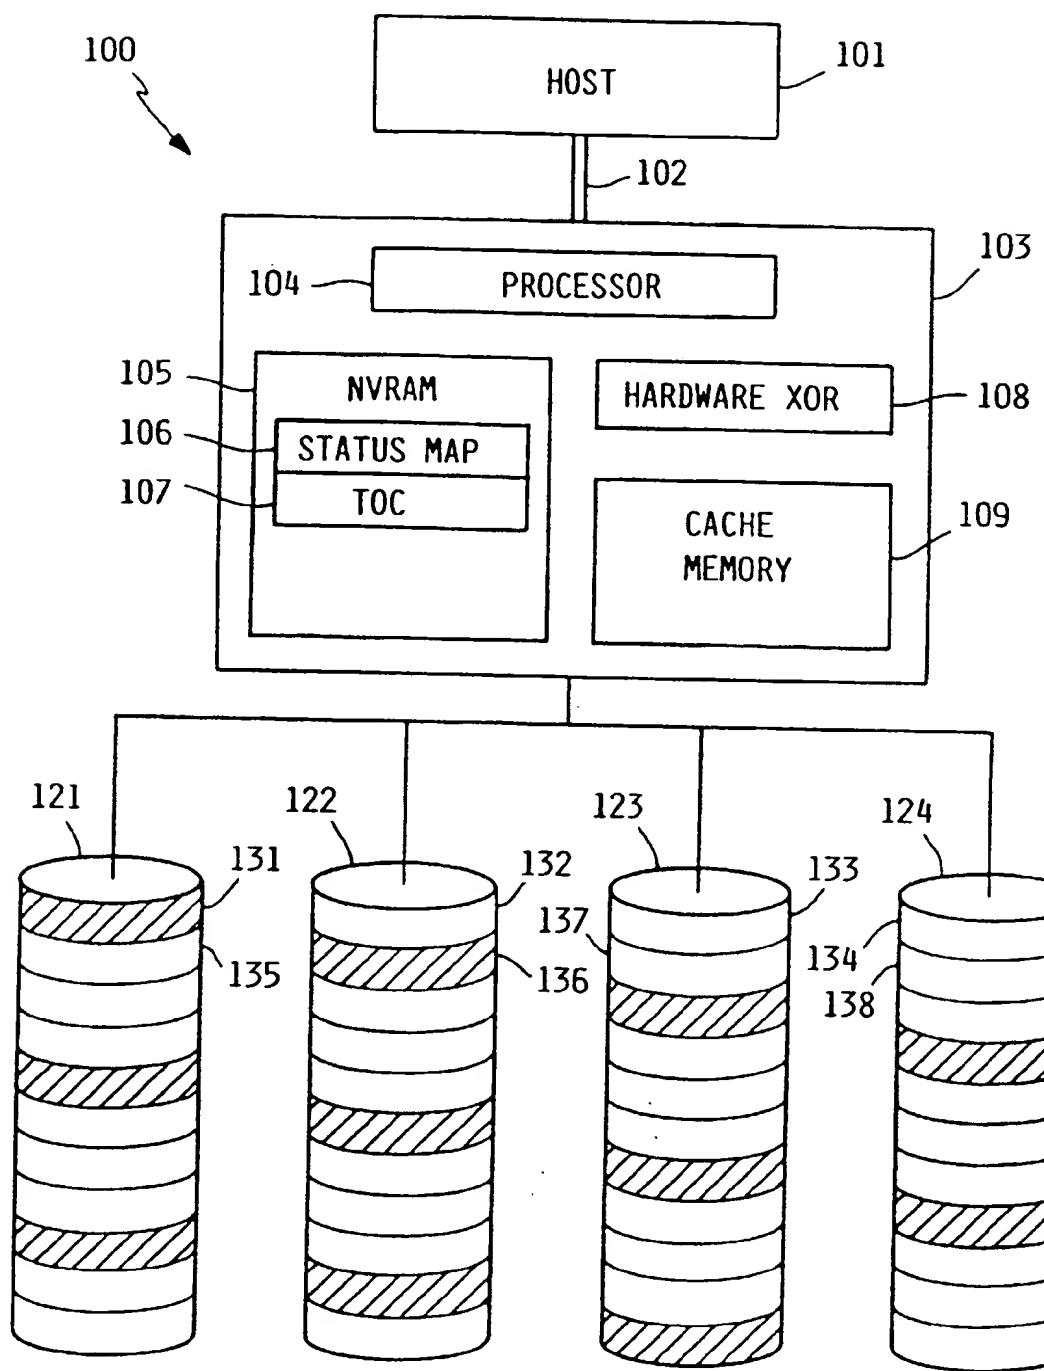


FIG. 1

BLOCK ADDRESS	STATUS	PARITY BLOCK ADDRESS
131	1	131
135	0	136
⋮	⋮	⋮
132	1	131
136	0	136
⋮	⋮	⋮
133	1	131
137	0	136
⋮	⋮	⋮

FIG. 2

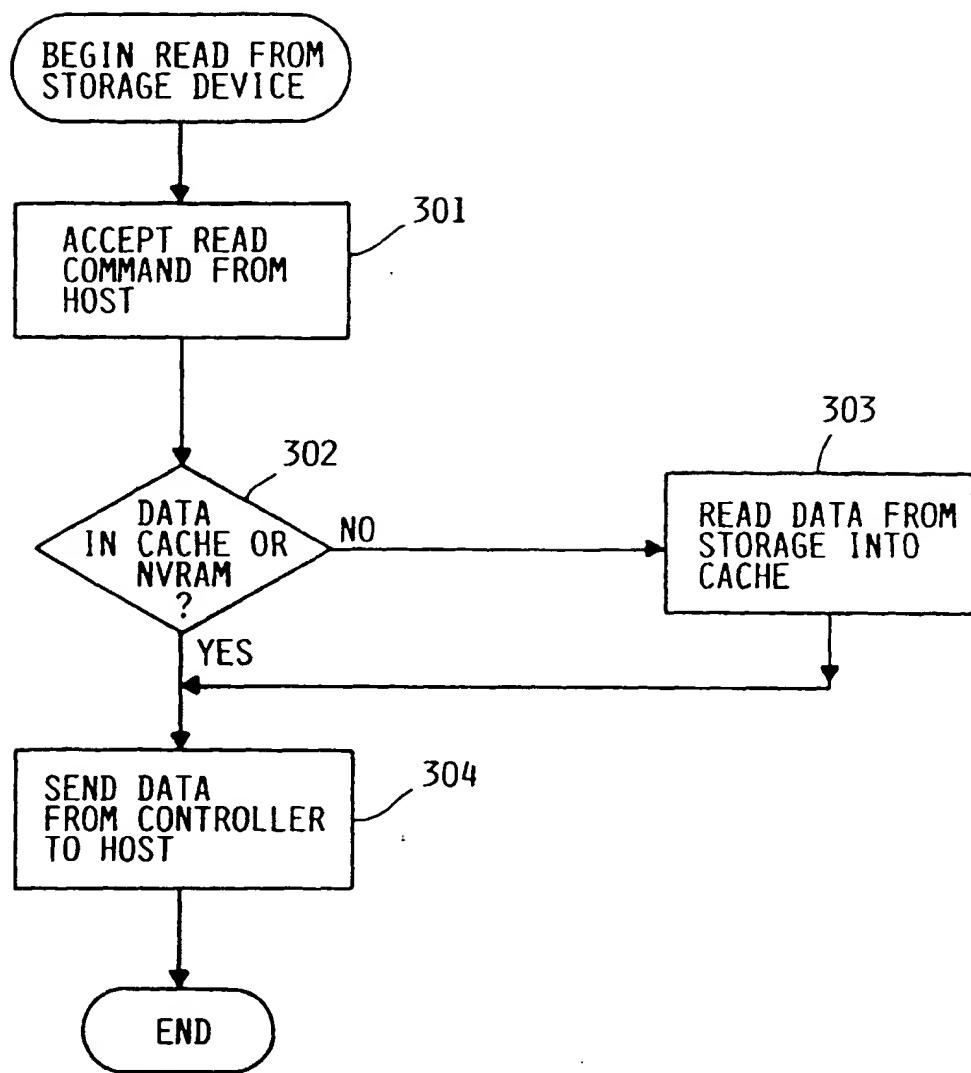


FIG. 3

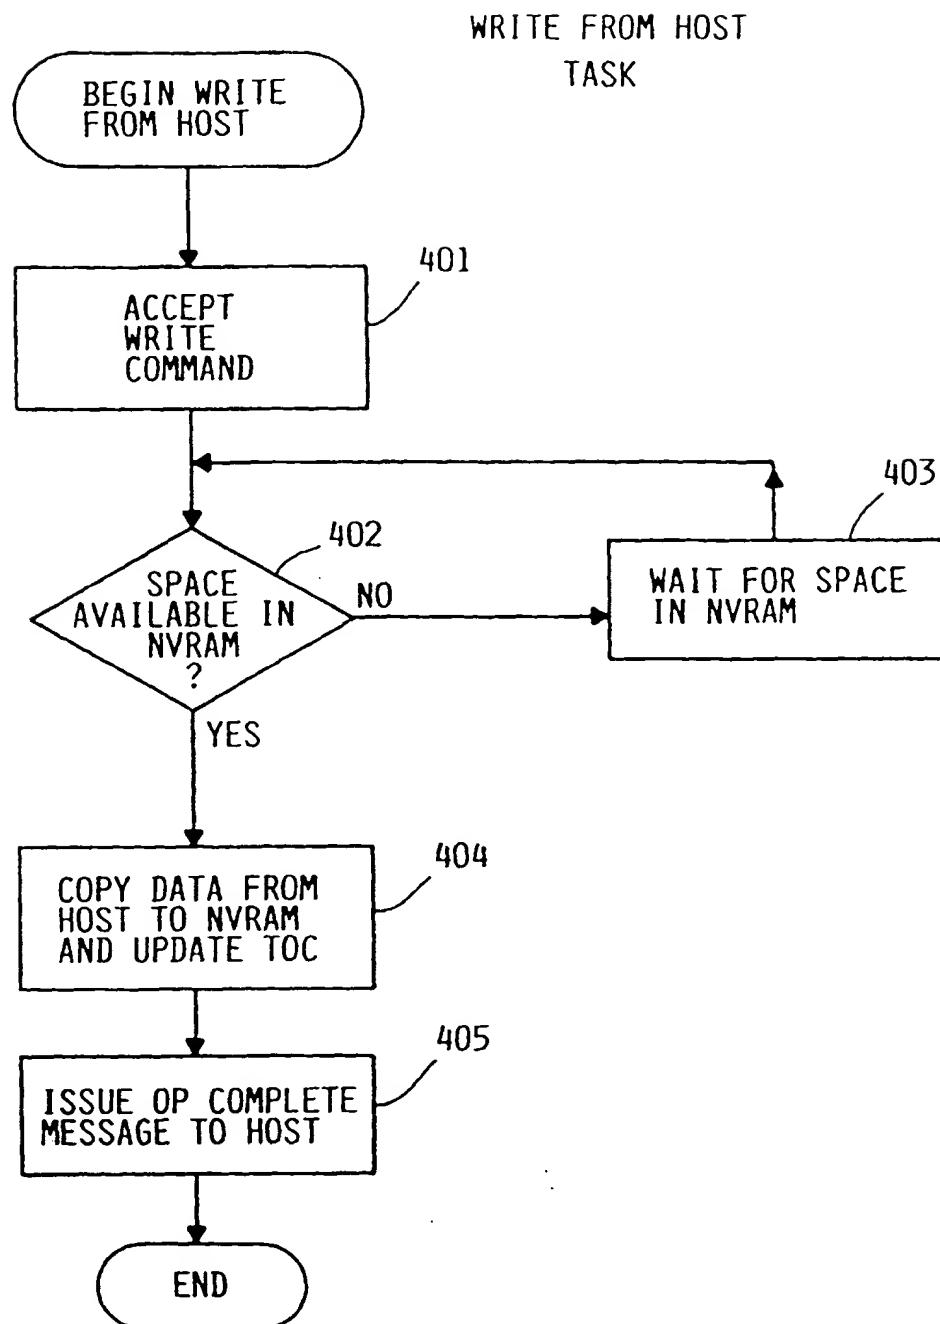


FIG. 4

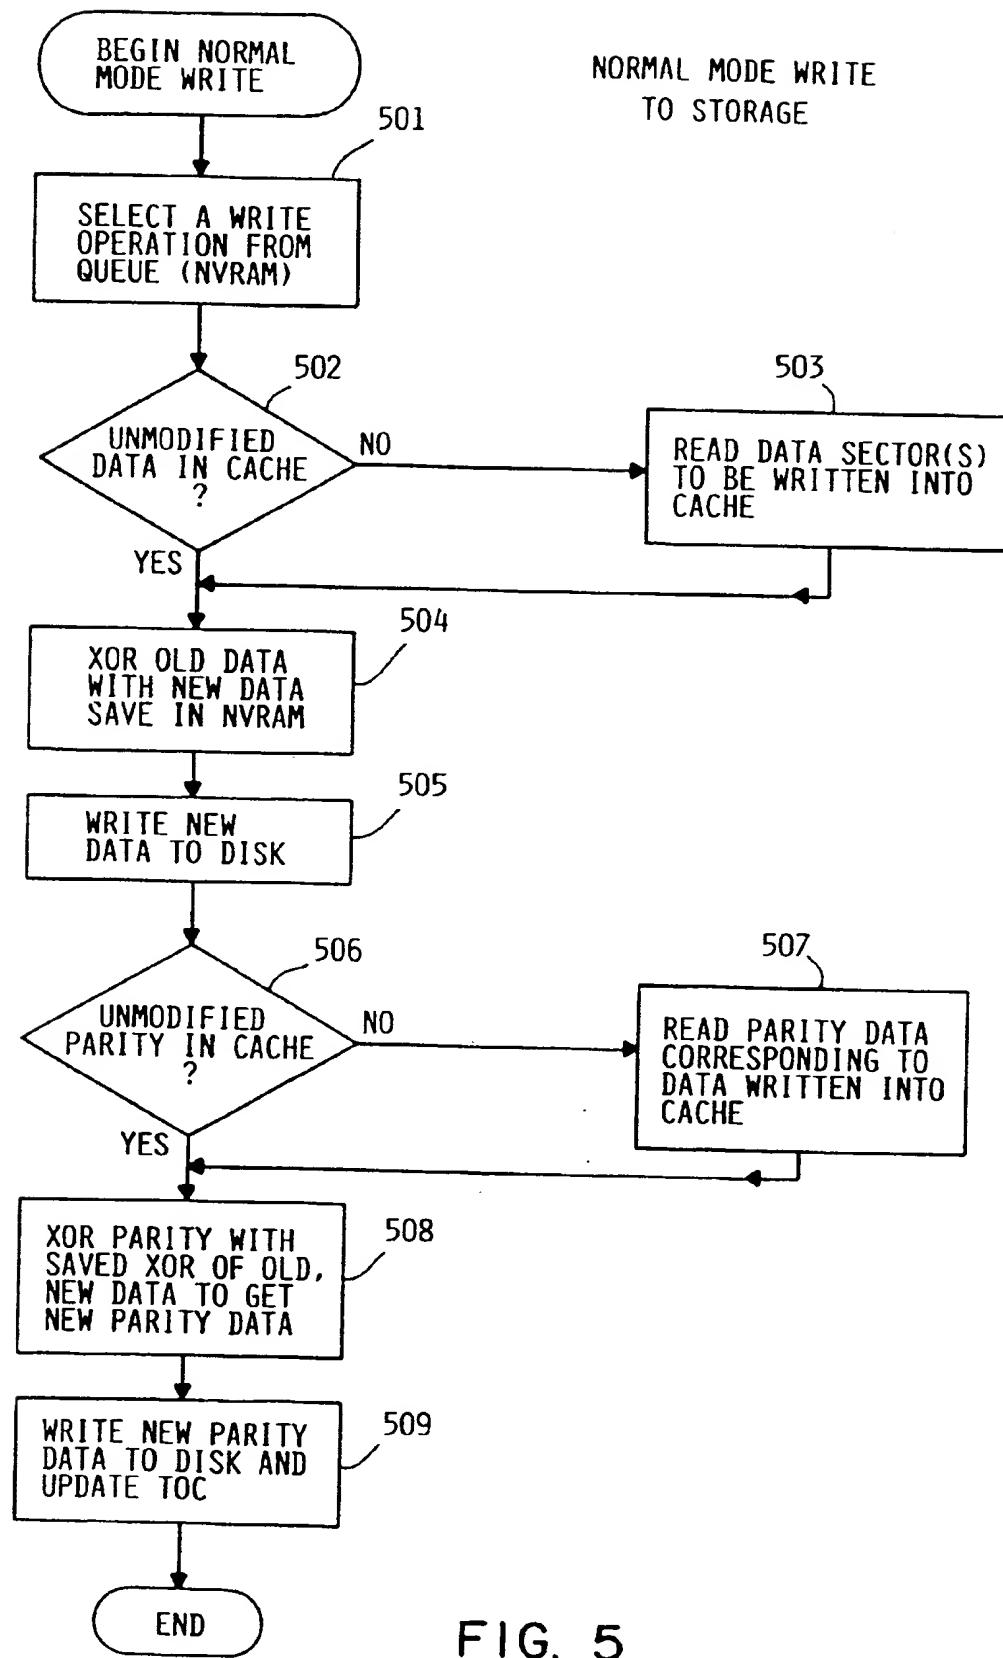


FIG. 5

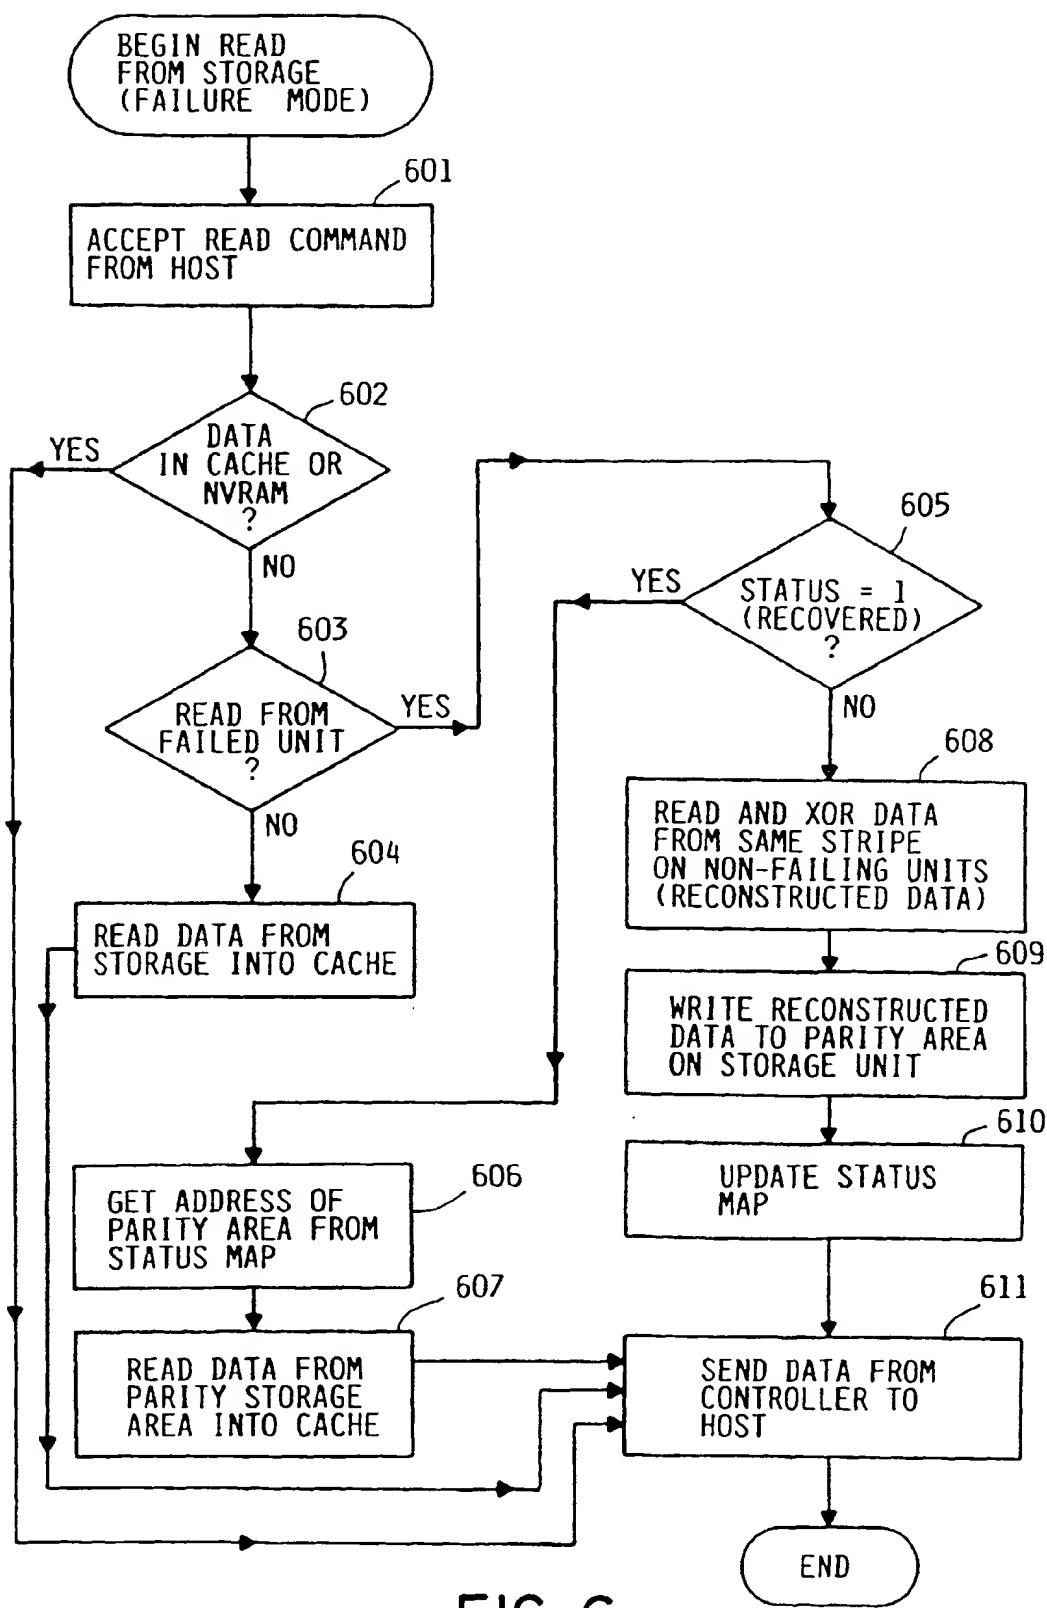


FIG. 6

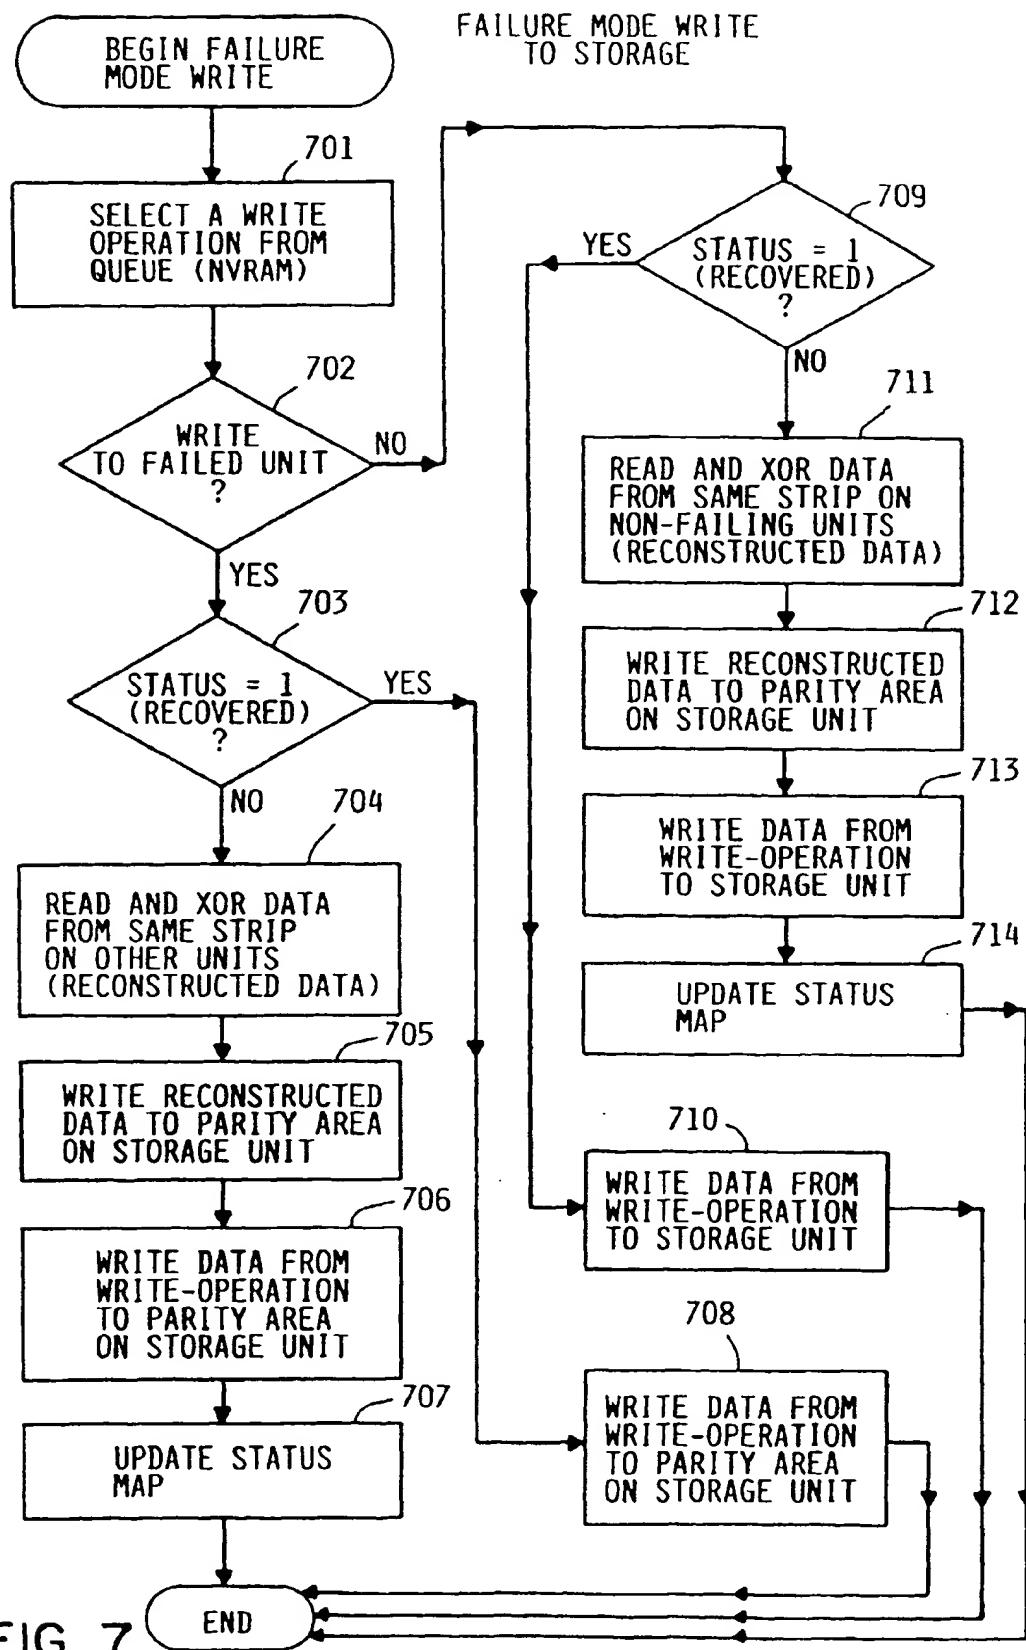


FIG. 7

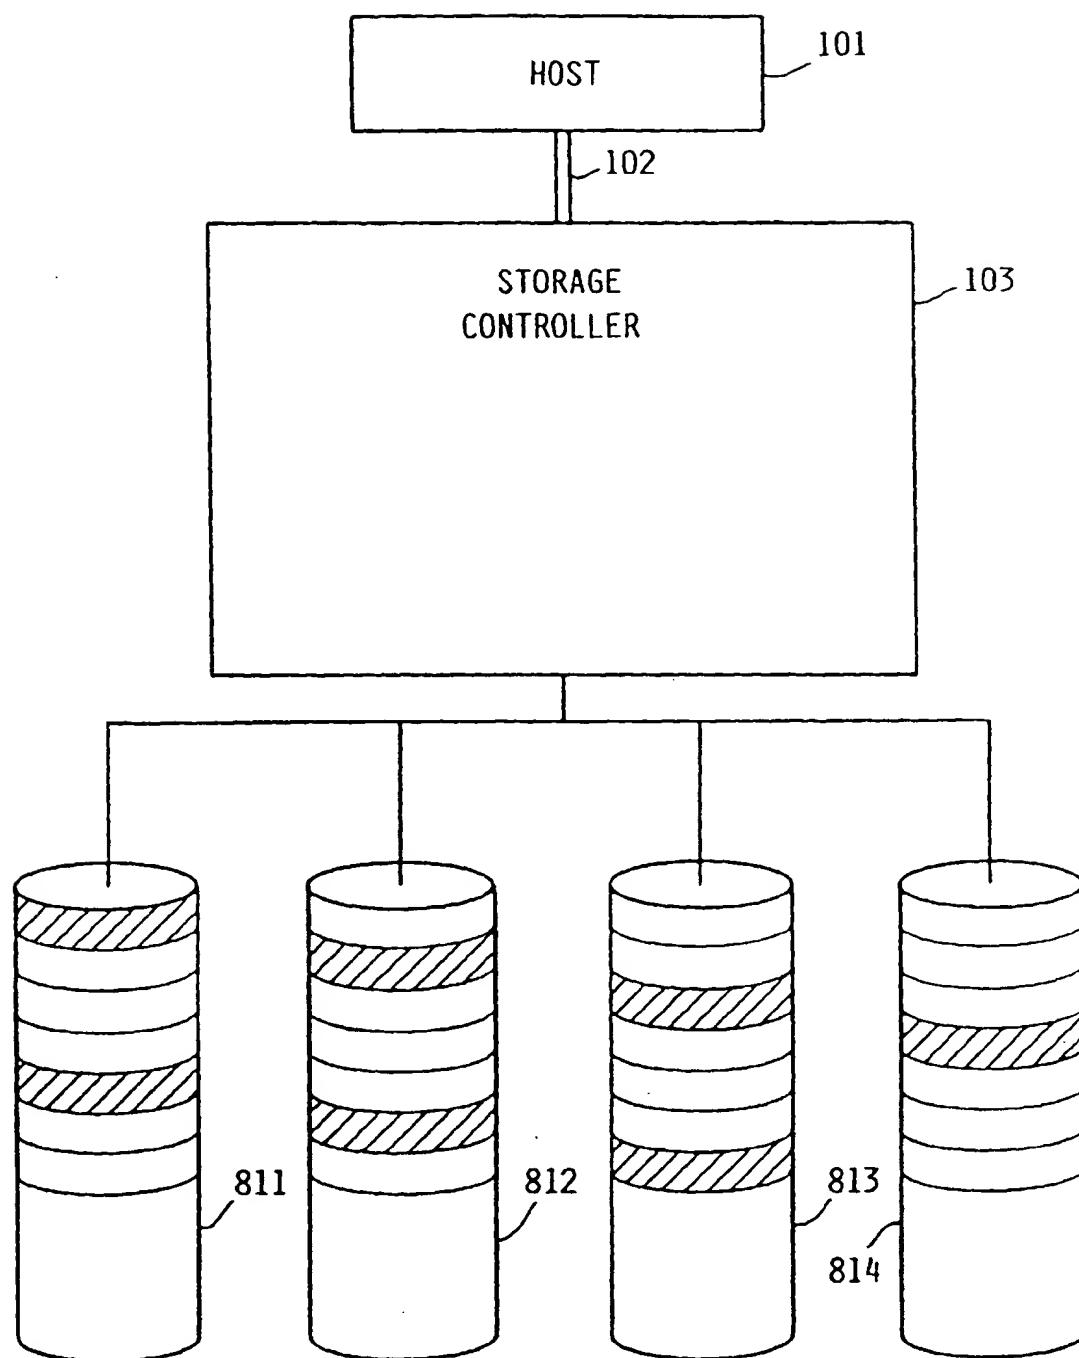


FIG. 8